

Better Matching, Less Forgetting: A Quality-Guided Matcher for Transformer-based Incremental Object Detection

Qirui Wu,¹ Shizhou Zhang,^{1*} De Cheng,^{2,1*} Yinghui Xing,¹ Lingyan Ran,¹ Dahu Shi,^{3,4} Peng Wang¹

¹ASGO, School of Computer Science, Northwestern Polytechnical University, China

²School of Telecommunications Engineering, Xidian University, China

³CCAI, Zhejiang University, China

⁴Hikrobot Co., Ltd.

wuqirui@mail.nwpu.edu.cn, szzhang@nwpu.edu.cn, dcheng@xidian.edu.cn, xyh_7491@nwpu.edu.cn, lran@nwpu.edu.cn, shidahu@zju.edu.cn, peng.wang@nwpu.edu.cn

Abstract


Incremental Object Detection (IOD) aims to continuously learn new object classes without forgetting previously learned ones. A persistent challenge is catastrophic forgetting, primarily attributed to background shift in conventional detectors. While pseudo-labeling mitigates this in dense detectors, we identify a novel, distinct source of forgetting specific to DETR-like architectures: background foregrounding. This arises from the exhaustiveness constraint of the Hungarian matcher, which forcibly assigns every ground truth target to one prediction, even when predictions primarily cover background regions (i.e., low IoU). This erroneous supervision compels the model to misclassify background features as specific foreground classes, disrupting learned representations and accelerating forgetting. To address this, we propose a Quality-guided Min-Cost Max-Flow (Q-MCMF) matcher. To avoid forced assignments, Q-MCMF builds a flow graph and prunes implausible matches based on geometric quality. It then optimizes for the final matching that minimizes cost and maximizes valid assignments. This strategy eliminates harmful supervision from background foregrounding while maximizing foreground learning signals. Extensive experiments on the COCO dataset under various incremental settings demonstrate that our method consistently outperforms existing state-of-the-art approaches.

Code — <https://github.com/fanrena/Q-MCMF>

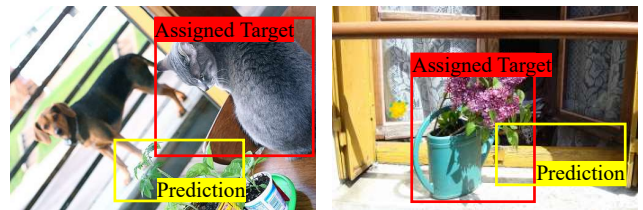
1 Introduction

Incremental Object Detection (IOD) aims to enable models to learn new object classes continuously while preserving performance on previously learned classes. A critical challenge in IOD is catastrophic forgetting, where conventional detectors (e.g., Faster R-CNN (Ren et al. 2015), DETR (Carion et al. 2020)) tend to lose previously acquired knowledge. Although significant efforts have been devoted to mitigating catastrophic forgetting in object detectors, it remains a persistent and challenging problem in computer vision.

*Shizhou Zhang and De Cheng are co-corresponding authors. This work is done during Qirui Wu’s internship at Hikrobot. Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Proportion of Matches with IoU < Threshold



(b) Examples of Background Foregrounding

Figure 1: Quantitative evidence and examples of background foregrounding: (a) shows the proportion of matches with IoU < IoU threshold of Base/New classes at the 40th epoch of 70-10 second phase. (b) shows two examples of background foregrounding.

Previous research identifies background shift (Cermelli et al. 2020) as a major source of forgetting. This occurs when objects belonging to previously learned classes (e.g., “airplane”) appear in the training data of a new task but are incorrectly labeled as “background”. Optimizing the detector using these mislabeled foreground objects as background not only disrupt the learned foreground representations but also corrupt the background distribution, leading to significantly more severe catastrophic forgetting than inter-class confusion among foreground objects. Pseudo-labeling has emerged as a simple yet effective strategy to address background shift, applicable to detectors employing both dense predictions (Mo et al. 2024; Wu et al. 2025) and sparse predictions (Liu et al. 2023b; Zhang et al. 2025b).

Unlike dense detectors, DETR (Carion et al. 2020) distinguishes itself by performing end-to-end object detection

without requiring post-processing such as Non-Maximum Suppression (NMS). This capability primarily stems from its use of sparse object queries and the Hungarian matching algorithm. Crucially, the Hungarian matcher serves as the label assignment mechanism, enforcing a strict one-to-one correspondence between predictions and ground truth objects. A key characteristic of this matcher is that it exhaustively assigns each target to one prediction, even if the prediction and the target are not geometrically plausible. This exhaustive assignment property introduces a novel source of forgetting specific to DETR-like architectures, which we define as *background foregrounding*.

Background foregrounding arises when a prediction, primarily covering background regions, is forcibly matched to a target due to the exhaustiveness constraint of the Hungarian matcher (see Fig. 1 (b)). Consequently, the model is incorrectly supervised to classify background features as belonging to a specific foreground class. Similar to background shift, optimizing the model with this erroneous supervision disrupts the learned feature distributions for both the actual foreground classes and the background class, significantly contributing to catastrophic forgetting. Fig. 1 (a) shows the proportion of matches with IoU below the threshold. At 40th epoch, 5.7% (old) and 37.6% (new) of matches have $\text{IoU} < 0.7$, indicating substantial background has been assigned as foreground, degrading the detector’s plasticity and stability.

To address this unique challenge inherent in DETR-based IOD, we propose a **Quality-guided Min-Cost Max-Flow (Q-MCMF)** matcher. First, we construct a graph between predictions and ground truth objects with costs, flows, and quality (e.g. IoU) as the property of each edge, then prune all edges with IoU below a threshold to eliminate geometrically implausible matches. This yields a sparse graph where connections represent only valid foreground-background relationships. On this refined graph, we solve the matching problem with dual objectives: 1) Minimizing the total assignment cost to prioritize low-cost matches. 2) Maximizing the number of matched pairs to preserve critical supervision signals. Crucially, Q-MCMF does not force low-IoU matches, allowing background regions and foreground targets to stay unmatched. This eliminates erroneous supervision that drives background foregrounding. Simultaneously, the matcher maximizes valid matches to prevent slow convergence and performance drops due to sparse supervision after pruning, balancing quantity and quality for robust incremental supervision. Experiments on the COCO dataset in various settings shows the effectiveness of our method.

Our contributions can be summarized as follows:

- To the best of our knowledge, we are the first to identify that background foregrounding as a critical problem that causes catastrophic forgetting in incremental DETR.
- To address the background foregrounding problem, we propose a generic Quality guided MCMF Matcher that eliminates low-quality matching while establishing more positive matches.
- The proposed approach outperforms existing methods across a variety of single-step and multi-step settings, highlighting its significant effectiveness.

2 Related Works

2.1 Detection Transformers

DETR (Carion et al. 2020) pioneered end-to-end object detection by eliminating hand-crafted components such as NMS through its use of sparse object queries and a one-to-one Hungarian matcher. While this paradigm simplifies the detection pipeline, DETR encounters significant challenges: the implicit semantics of its object queries (Meng et al. 2021; Liu et al. 2022; Xing et al. 2024) and slow training convergence (Zhu et al. 2020; Li et al. 2022; Chen et al. 2023; Jia et al. 2023; Huang et al. 2025). Addressing the convergence issue specifically, sparse positive supervision has been identified as a primary bottleneck. These strategies (Chen et al. 2023; Jia et al. 2023; Huang et al. 2025) demonstrate substantial improvements in accelerating convergence speed and boosting performance.

While existing methods effectively increase the quantity of supervision, we contend that for incremental learning scenarios, the quality of positive matches is equally critical.

2.2 Incremental Object Detection

Incremental Object Detection (IOD) faces the dual challenge of preserving previously acquired detection capabilities while assimilating novel object classes. This complexity stems from two primary vulnerabilities: classification/localization forgetting (Wu et al. 2025; Zhang et al. 2025a) and the background shift phenomenon (Liu et al. 2023a; Mo et al. 2024; Zhang et al. 2025c; Kim et al. 2024). Background shift (Cermelli et al. 2020) occurs when instances of previously learned classes in new-task training data are mislabeled as background, progressively corrupting established representations. Common mitigation strategies include pseudo-labeling (Mo et al. 2024; Liu et al. 2023b; Luo et al. 2025) and exemplar replay (Liu et al. 2023b,a; Kim et al. 2024; Zhang et al. 2025c; Wu et al. 2025), which aim to preserve knowledge of prior classes.

Diverging from these approaches, our work investigates a more fundamental aspect: the role of label assignment mechanisms within DETR-based IOD frameworks, where we identify critical limitations requiring novel solutions.

2.3 Label Assignment

Defining positive/negative training samples is fundamental in object detection, as the quality of assignment critically impacts performance. Current assignment strategies operate under predefined correspondence rules, broadly categorized into two paradigms: 1. One-to-many assigners (e.g., Max-IoU, ATSS (Zhang et al. 2020), OTA (Ge et al. 2021a), SimOTA (Ge et al. 2021b)) allocate multiple predictions per ground truth. 2. One-to-one assigners enforce strict pairwise matching, exemplified by DETR’s (Carion et al. 2020) Hungarian algorithm, enabling NMS-free end-to-end detection.

Our analysis identifies a flaw in the Hungarian matcher: its exhaustive matching forces low-quality predictions to targets, causing background foregrounding. This disrupts feature distributions catastrophically during incremental learning via erroneous supervision. Hence, we propose Q-MCMF Matcher to resolve this problem.

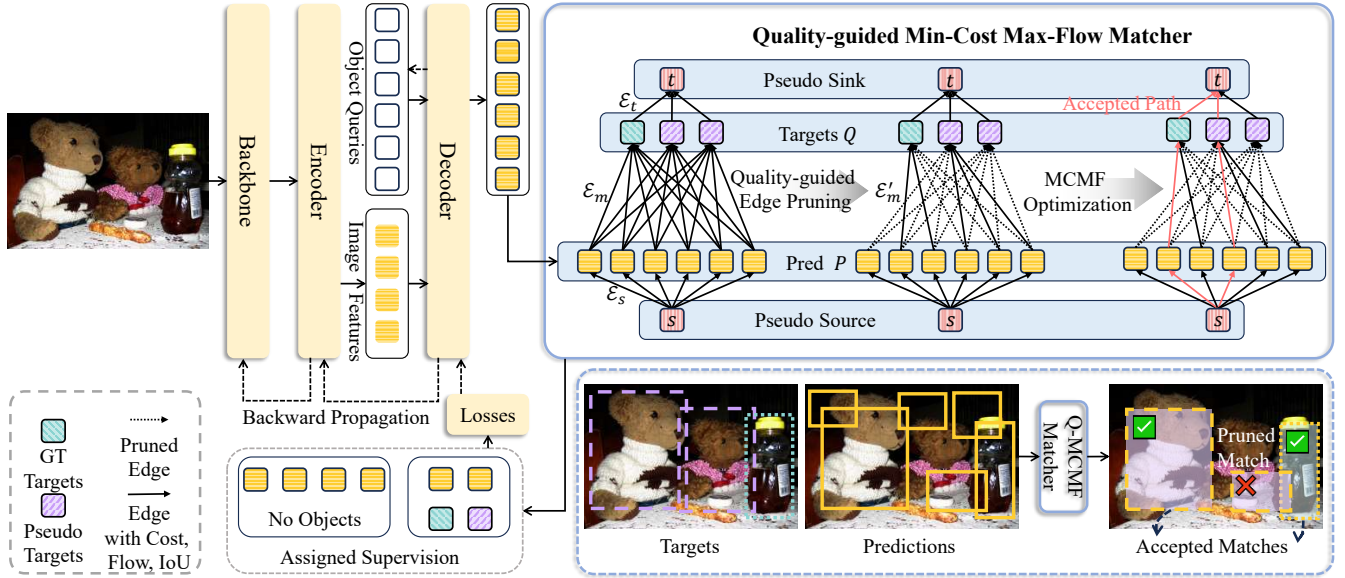


Figure 2: Overall training pipeline of Deformable DETR with Q-MCMF matcher.

3 Method

3.1 Preliminary

IOD Formulation. In IOD, model training is structured across n sequential learning stages, where each stage introduces a novel set of classes for detection. Formally, let $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n\}$ represent the complete class set incrementally acquired by detector \mathcal{F} , with $\mathcal{T}_i \cap \mathcal{T}_j = \emptyset$ for all $i \neq j$. The training dataset for stage k is denoted as $\mathcal{D}_k = \{\mathcal{X}_k, \mathcal{Y}_k\}$, where \mathcal{X}_k contains input images and \mathcal{Y}_k provides corresponding annotations. A critical characteristic of IOD is that while images in \mathcal{X}_k may contain objects from any class in \mathcal{T} , only those belonging to \mathcal{T}_k are explicitly annotated. The primary challenge lies in updating the detector from \mathcal{F}_{k-1} to \mathcal{F}_k using only the current dataset \mathcal{D}_k (i.e., without access to previous data $\{\mathcal{D}_1, \dots, \mathcal{D}_{k-1}\}$), so as to improve its capability on \mathcal{T}_k while avoiding catastrophic forgetting on all previous classes $\bigcup_{i=1}^{k-1} \mathcal{T}_i$.

Detection Transformer. Our work focuses on the DETR architecture, specifically Deformable DETR (Zhu et al. 2020). Deformable DETR is a fully end-to-end object detector comprising four key components: 1) a backbone that extracts initial multi-scale feature maps from the input image; 2) a transformer encoder that encodes these features; 3) object queries, a set of learnable embeddings encoding positional and semantic priors, that interact with the encoded features to locate objects; and 4) a transformer decoder that uses the object queries to attend to relevant regions within the encoder’s output feature map. Through iterative cross-attention layers, the decoder refines each query into an object representation. These final decoded representations are fed into prediction heads for bounding box regression and category classification.

The source of DETR’s end-to-end characteristic lies in its prediction-target matcher. Predictions are assigned to

ground truth targets via the Hungarian Matcher, which finds a strict one-to-one, bipartite assignment between predictions and targets by minimizing global cost. This matching strategy eliminates the need for NMS post-processing, enabling DETR’s end-to-end nature.

However, the exhaustive matching property of the Hungarian algorithm—which compels every ground truth target to be assigned a prediction, even when no geometrically compatible prediction exists—induces knowledge erosion. This occurs through misclassification of background as foreground, disrupting the model’s stability-background balance.

Hungarian Matcher. In DETR, the Hungarian matcher serves as the label assigner to assign labels to predictions in a strict one-to-one correspondence. Specifically, let the predictions be $P = \{p_i \mid i \in \mathbb{N}, 1 \leq i \leq N_p\}$ where N_p is the number of predictions, and the targets be $Q = \{q_i \mid i \in \mathbb{N}, 1 \leq i \leq N_q\}$ where N_q is the number of targets. There is a complete bipartite graph $\mathcal{G}_m = (P, Q; \mathcal{E}_m)$ where $\mathcal{E}_m = \{(p_i, q_j) \mid \forall p_i \in P, \forall q_j \in Q\}$ is the collection of edges in the graph with a non-negative cost $c(p_i, q_j)$, usually to be the loss between p_i and q_j . The $\mathcal{C} = \{c(p_i, q_j) \mid \forall p_i \in P, \forall q_j \in Q\}$ is also known as the cost matrix. The goal of the Hungarian matcher is to find an assignment of the targets to the predictions $\mathcal{M} \subset \mathcal{E}_m$ while the total cost of assignment is minimum. That is to say, the optimization goal of the Hungarian matcher is

$$\begin{aligned} & \arg \min_{\mathcal{M}} \sum_{(p_i, q_j) \in \mathcal{M}} c(p_i, q_j), \\ \text{s.t.} & \begin{cases} \sum_{j=1}^{N_p} \mathbb{I}[(p_i, q_j) \in \mathcal{M}] \leq 1, \forall p_i \in P \\ \sum_{i=1}^{N_q} \mathbb{I}[(p_i, q_j) \in \mathcal{M}] \leq 1, \forall q_j \in Q \\ |\mathcal{M}| = \min(N_p, N_q), \end{cases} \quad (1) \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. Here, the first two restrictions limit edges in \mathcal{M} from sharing vertex, ensuring a

strict one-to-one assignment. The last constraint dictates that the cardinality of \mathcal{M} equals the smaller of the cardinalities of P and Q . This ensures that every single node in P or Q is exhaustively assigned to a target/prediction.

The fundamental limitation lies in the complete bipartite graph structure and its requirement for perfect matching on one node set as shown in Eq. 1. The Hungarian matcher’s exhaustive assignment induces background foregrounding, which catastrophically corrupts feature distributions in incremental learning scenarios. To address this, we propose a simple yet effective Quality-guided Min-cost Max-flow (Q-MCMF) Matcher.

3.2 Quality-guided Min-Cost Max-Flow Matcher

To mitigate the background foregrounding problem, our matching framework should satisfy four key properties: 1) Strict one-to-one correspondence in assignments. 2) Minimized matching cost over the selected pairs. 3) Selective exclusion capability that allows targets without adequate matches to be excluded from the assignment, thereby eliminating background foregrounding. 4) Maximized matched pairs, as sparse matching hampers model from convergence and limits performance as suggested in prior works (Chen et al. 2023; Jia et al. 2023; Huang et al. 2025). To meet these requirements, we reformulate the matching task as a min-cost max-flow (MCMF) problem. This framework provides flexible cost matrix manipulation while simultaneously satisfying all four requirements through a unified optimization process.

Graph Construction. To start, we transform the complete bipartite graph \mathcal{G}_m into a directed flow network $\mathcal{G} = (s, t, P, Q; \mathcal{E})$, where s and t are pseudo source and pseudo sink respectively. As shown in Fig. 2, the directed edge set \mathcal{E} is defined as the union of three edge sets $\mathcal{E} = \mathcal{E}_s \cup \mathcal{E}_m \cup \mathcal{E}_t$ where

$$\begin{cases} \mathcal{E}_s = \{(s, p_i) \mid p_i \in P\} \\ \mathcal{E}_m = \{(p_i, q_j) \mid p_i \in P, q_j \in Q\} \\ \mathcal{E}_t = \{(q_j, t) \mid q_j \in Q\}. \end{cases} \quad (2)$$

Every edge has three attributes: cost $c(p_i, q_j)$, flow $f(p_i, q_j)$, and IoU $\phi(p_i, q_j)$ between p_i and q_j . For every $p_i \in P$ and $q_j \in Q$, we assign $c(s, p_i) = 0$, $c(q_j, t) = 0$, $\phi(s, p_i) = 1$, and $\phi(q_j, t) = 1$ to ensure zero contribution of \mathcal{E}_s and \mathcal{E}_t to the results. The flow capacity is set to $f(e) = 1$ for every edge $e \in \mathcal{E}$ to ensure one-to-one assignment.

Quality-guided Edge Pruning. With P and Q densely connected by \mathcal{E}_m , solving the min-cost max-flow problem on \mathcal{G} deteriorates to a matching problem on \mathcal{G}_m , and it only fosters the minimization of cost. To reduce background foregrounding, we introduce quality-guided edge pruning to prune edges with low IoU. Specifically, we refine \mathcal{E}_m by

$$\mathcal{E}'_m = \{(p_i, q_j^{old}) \mid \phi(p_i, q_j^{old}) \geq \alpha\} \cup \{(p_i, q_j^{new}) \mid \phi(p_i, q_j^{new}) \geq \beta\}, \quad (3)$$

where q_j^{old} and q_j^{new} represent target labels sourced from pseudo-labels (learned tasks) or ground truth (current task), respectively. α and β is the different threshold for different targets.

(1)	Methods	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
40-40	LwF*	17.2	25.4	18.6	7.9	18.4	24.3
	CL-DETR*	42.0	60.1	45.9	24.0	45.3	55.6
	SDDGR*	<u>43.0</u>	<u>62.1</u>	<u>47.1</u>	<u>24.9</u>	<u>46.9</u>	<u>57.0</u>
	DyQ-DETR*	42.4	60.5	45.9	23.9	46.3	56.7
	RILOD	29.9	45.0	32.0	15.8	33.0	40.5
	SID	34.0	51.4	36.3	18.4	38.4	44.9
	ERD	36.9	54.5	39.6	21.3	40.4	47.5
	CL-DETR	39.2	56.1	42.6	21.0	42.8	52.6
	DyQ-DETR	41.4	59.7	44.9	24.1	45.2	54.3
	DCA	42.8	58.4	–	–	–	–
	Ours	44.3	59.5	49.2	27.1	49.9	59.8
	70-10	LwF*	7.1	12.4	7.0	4.8	9.5
CL-DETR*		40.4	58.0	43.9	23.8	43.6	53.5
SDDGR*		40.9	59.5	44.8	23.9	44.7	54.0
DyQ-DETR*		<u>42.4</u>	<u>60.4</u>	<u>46.3</u>	<u>24.5</u>	<u>45.7</u>	<u>57.5</u>
RILOD		24.5	37.9	25.7	14.2	27.4	33.5
SID		32.8	49.0	35.0	17.1	36.9	44.5
ERD		34.9	51.9	37.4	18.7	38.8	45.5
CL-DETR		35.8	53.5	39.5	19.4	41.5	46.1
DyQ-DETR		39.5	56.4	43.1	22.5	43.1	53.0
DCA		41.3	59.2	–	–	–	–
Ours		43.4	62.0	47.6	28.2	46.9	56.0

Table 1: IOD results (%) on COCO 2017 under the 40-40 and 70-10 setting using protocol (1). * indicates results with exemplar replay. The best performance in each is presented in bold, and the second best is presented with underline.

The rationale for employing distinct thresholds for different targets stems from the disparate predictive capabilities of the model with respect to seen and unseen classes. For seen classes, accurate predictions and effective assignment are more readily achievable. Conversely, for unseen classes, achieving high-quality assignment proves challenging due to the model’s lack of prior exposure. A high threshold can reduce the number of positive samples and slow down the convergence rate, whereas a low threshold, although it increases the positive matches, intensifies the background foregrounding issue. Consequently, the parameters α and β govern the balance between stability and plasticity within our method.

MCMF Optimization. The objective is to find a set of paths $\mathcal{P} = \{s \rightarrow p_i \rightarrow q_j \rightarrow t \mid p_i \in P, q_j \in Q\}$ that simultaneously minimizes total cost and maximizes flow. Following quality-guided edge pruning, we formulate the MCMF problem’s objectives as:

$$\begin{aligned} & \arg \min_{\mathcal{P}} \sum_{(p_i, q_j) \in \mathcal{P}} c(p_i, q_j), \\ & \text{s.t.} \begin{cases} \sum_{j=1}^{N_p} \mathbb{I}[(p_i, q_j) \in \mathcal{P}] \leq 1, \forall p_i \in P \\ \sum_{i=1}^{N_q} \mathbb{I}[(p_i, q_j) \in \mathcal{P}] \leq 1, \forall q_j \in Q \\ |\mathcal{P}| = \arg \max_{|\mathcal{P}|} \left(\sum_{(s, p_i, q_j, t) \in \mathcal{P}} f(s, p_i, q_j, t) \right). \end{cases} \end{aligned} \quad (4)$$

Here, akin to Eq. 1, the first two restrictions ensure that paths in \mathcal{P} do not share intermediate edges and guarantee strict

(2)	Methods	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
40-40	LwF*	23.9	41.5	25.0	12.0	26.4	33.0
	CL-DETR*	37.5	55.1	40.3	20.9	40.8	50.7
	DyQ-DETR*	<u>39.7</u>	<u>57.5</u>	<u>43.0</u>	<u>21.6</u>	<u>42.9</u>	<u>53.8</u>
	iCaRL	33.4	52.0	36.0	18.0	36.4	45.5
	ERD	36.0	55.2	38.7	19.5	38.7	49.0
	CL-DETR	36.2	52.6	39.5	18.7	39.5	49.4
	DyQ-DETR	39.1	57.1	42.5	21.3	42.7	51.8
	Ours	42.4	60.5	46.4	25.5	46.5	55.2
70-10	LwF*	24.5	36.6	26.7	12.4	28.2	35.2
	CL-DETR*	40.1	57.8	43.7	23.2	43.2	52.1
	DyQ-DETR*	<u>41.9</u>	<u>60.1</u>	<u>45.8</u>	<u>24.1</u>	<u>45.3</u>	<u>55.8</u>
	iCaRL	35.9	52.5	39.2	19.1	39.4	48.6
	ERD	36.9	55.7	40.1	21.4	39.6	48.7
	CL-DETR	34.0	48.0	37.2	15.5	37.7	49.7
	DyQ-DETR	39.6	57.6	43.5	23.4	43.3	51.8
	Ours	43.1	61.3	47.5	27.9	46.7	55.7

Table 2: IOD results (%) on COCO 2017 under the 40-40 and 70-10 setting using protocol (2). * indicates results with exemplar replay. The best performance in each is presented in bold, and the second best is presented with underline.

one-to-one matching. The last constraint limits the flow of path set \mathcal{P} to the maximum flow in graph \mathcal{G} . This not only ensures the number of maximum matching pairs but also allows implausible matches to be discarded in the final assignment, thereby preventing the background foregrounding problem. The solution reduces to the intermediate edges $\mathcal{M}' = \{(p_i, q_j) \mid \forall (s, p_i, q_j, t) \in \mathcal{P}\}$, which constitute the final matching assignment.

This optimization naturally yields an assignment of the prediction target \mathcal{M}' that satisfies our four design criteria and effectively mitigates the background foregrounding problem.

3.3 Training Objectives

After obtaining the prediction-target assignment \mathcal{M}' , we compute the same loss function as that used in conventional Deformable DETR. Here, we denote p_i^c and p_i^b as the class prediction and its bounding box prediction of p_i while q_j^c and q_j^b as the class label and its bounding box annotation of q_i . The loss function for each matched pair $(p_i, q_j) \in \mathcal{M}'$ consists of

$$\mathcal{L}_{fg} = \sum_{(p_i, q_j) \in \mathcal{M}'} (\lambda_{\text{focal}} \mathcal{L}_{\text{focal}}(p_i^c, q_j^c) + \lambda_{L1} \mathcal{L}_{L1}(p_i^b, q_j^b) + \lambda_{\text{giou}} \mathcal{L}_{\text{giou}}(p_i^b, q_j^b)), \quad (5)$$

where $\mathcal{L}_{\text{focal}}$ denotes the focal loss for predicted class p_i^c and target class q_j^c , \mathcal{L}_{L1} and $\mathcal{L}_{\text{giou}}$ represent bounding box regression losses between predicted box p_i^b and target box q_j^b , and λ terms are the loss coefficients to balance the trade-off between different loss components in the optimization objective. Predictions not assigned to any target ($P_u = \{p_i \mid p_i \notin \mathcal{M}'\}$) are supervised as background class. The classification loss for these predictions is the standard focal loss

Methods	40-10 \times 4				40-20 \times 2	
	<u>+40</u>	<u>-50</u>	<u>+50</u>	<u>-60</u>	<u>+60</u>	<u>-80</u>
CF	5.8	5.7	6.3	3.3	10.7	9.4
RILOD	25.4	11.2	10.5	8.4	27.8	15.8
SID	34.6	24.1	14.6	12.6	34.0	23.8
ERD	36.4	30.8	26.2	20.7	36.7	32.4
CL-DETR*	–	–	–	28.1	–	35.3
ACF	39.1	35.4	32.0	30.3	39.3	36.6
SDDGR*	42.3	40.6	40.0	36.8	42.5	41.1
DCA	44.0	<u>41.1</u>	<u>39.2</u>	<u>37.2</u>	<u>42.7</u>	<u>40.3</u>
Ours	<u>43.2</u>	41.2	40.0	38.0	44.3	42.3

Table 3: IOD results (%) on COCO 2017 under the 40-10 \times 4 and 40-20 \times 2 setting using protocol (1). * indicates the results are obtained with exemplar replay. The best performance in each is presented in bold, and the second best is presented with underline.

formulated as

$$\mathcal{L}_{\text{bg}} = \sum_{p_i \in P_u} \mathcal{L}_{\text{focal}}(p_i^c, \emptyset), \quad (6)$$

where \emptyset represents background. Finally, the overall loss function is computed as

$$\mathcal{L} = \mathcal{L}_{fg} + \lambda_{\text{bg}} \mathcal{L}_{\text{bg}}. \quad (7)$$

4 Experiments

4.1 Experimental Setup

Dataset and Evaluation Metrics. For fair comparison, we follow the training protocols as in previous works (Liu et al. 2023b; Shmelkov, Schmid, and Alahari 2017) and evaluate our method on the widely adopted COCO 2017 dataset, which consists of 80 different classes in natural scenes. The AP, AP₅₀, AP₇₅, AP_S, AP_M, and AP_L of standard COCO metrics are used for performance evaluation.

Protocols. We evaluated our methods with two different protocols. (1) We follow previous works (Shmelkov, Schmid, and Alahari 2017) and split dataset with classes. In this protocol, images for training are overlapped for different tasks while their annotations are not. (2) Following (Liu et al. 2023b), we split the dataset into non-overlapping splits for training. In this protocol, both classes and unseen data are introduced sequentially and learned incrementally, which is a more practical simulation to the real-world application. Therefore, we evaluate our method with protocol (2) unless specified. We evaluate our method with two-phase and multiple-phase settings. For two-phase setting (A - B), we evaluate with 70-10 setting and 40-40 setting, where A are the number of base classes and B represents new classes. For multiple-phase setting (C - $D \times E$), we evaluate our method with 40-10 \times 4 and 40-20 \times 2 settings, where C denotes the number of base classes, D denotes the number of subsequent new classes, and E denotes the number of new phases.

Implementation Details. Following (Liu et al. 2023b), we build our method on top of the Deformable DETR without iterative bounding box refinement and the two-stage variant. To address background shift problem, we adopt DKD proposed in CL-DETR. The backbone we adopted is ResNet-50 (He et al. 2016) pre-trained on ImageNet (Deng et al. 2009). We train our detector with AdamW (Loshchilov and Hutter 2017) for 50 epoch. We set our quality guided edge pruning hyperparameter to $\alpha = 0.7$ and $\beta = 0.5$, which is determined through a grid search. We also implemented our method on DN-DETR (Li et al. 2022) and DAB-DETR (Liu et al. 2022) with the same hyperparameter as we used in Deformable DETR without further grid search for optimal hyperparameters. More implementation details can be found in the appendix.

4.2 Comparison with SOTA Methods

Two-Phase Settings. We compare our method with existing approaches, including exemplar replay-based methods (LwF* (Li and Hoiem 2017), CL-DETR* (Liu et al. 2023b), SDDGR* (Kim et al. 2024), DyQ-DETR* (Zhang et al. 2025b)) and exemplar-free methods (RILOD (Li et al. 2019), SID (Peng et al. 2021), ERD (Feng, Wang, and Yuan 2022), CL-DETR (Liu et al. 2023b), DyQ-DETR (Zhang et al. 2025b)), along with the previous state-of-the-art method DCA (Zhang et al. 2025a). Notably, our approach operates without exemplar replay. As presented in Tab. 1 under protocol (1), our method consistently outperforms all comparison methods across different incremental learning settings by significant margins. Our method surpasses its baseline, CL-DETR, by 5.1% and 7.6% in 40-40 and 70-10 settings in AP metric. It also surpasses previous best results by 1.0% and 1.3% in 70-10 and 40-40 respectively. It is worth noting that previous best results replay generative data (SD-DGR*) or 10% labeled data from the dataset (DyQ-DETR*) while ours preserve knowledge in an example-free manner.

Further results under protocol (2) (Tab. 2) demonstrate that our method surpasses DyQ-DETR by 3.5% in the 70-10 setting and achieves a 3.3% gain in the 40-40 setting. When compared to DyQ-DETR* with exemplar replay, our method also shows superior performance with gains of 1.2% (70-10) and 2.7% (40-40). The consistent performance gains across both protocols validate our method’s superiority to different task configurations. It also highlights the background foregrounding problem as a critical problem in incremental transformer detector.

Multiple-Phase Settings. In Tab. 3, we report the results under the more demanding multi-phase incremental learning setting, using the same data splits as prior works for fair comparison. Notably, in the 40-10 \times 4 configuration, our method outperforms exemplar-based methods CL-DETR and SDDGR by 9.9% and 1.2% respectively, while surpassing the previous state-of-the-art DCA by 0.8%. Similarly, in the 40-20 \times 2 setting, our approach achieves significant gains of 7.0% over CL-DETR, 1.2% over SDDGR, and 2.0% over DCA. These consistent improvements across different phase configurations demonstrate our method’s robustness in handling complex multi-phase learning scenarios. This also shows that solving background foregrounding

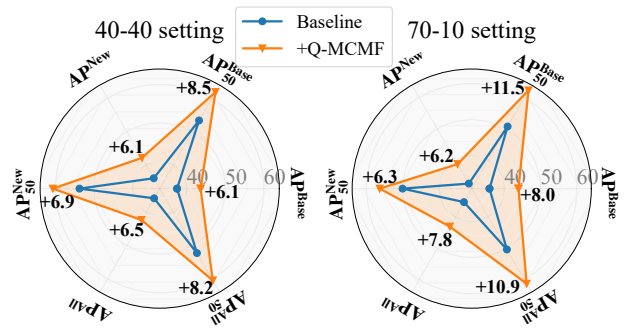


Figure 3: Effectiveness analysis of Q-MCMF Matcher.

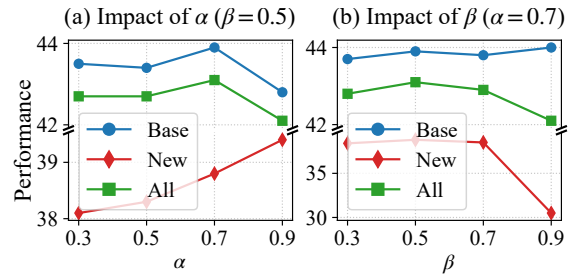


Figure 4: Stability-Plasticity Balance with different α and β .

problem can achieve consistent performance gain.

4.3 Ablation Study

Effectiveness Analysis of Q-MCMF Matcher. In Fig. 3, we present a comprehensive ablation study of our proposed Q-MCMF matcher, comparing its performance against the baseline. Our baseline implementation is CL-DETR without exemplar replay.

Regarding stability (Base classes), Q-MCMF demonstrates substantial improvements: +6.1% AP in the 40-40 setting and +8.0% AP in the 70-10 setting, indicating superior retention of prior knowledge. For plasticity (New classes), it achieves gains of +6.1% AP (40-40) and +6.2% AP (70-10), confirming effective adaptation to novel categories. The overall performance (All classes) shows consistent enhancement with improvements of +6.5% AP (40-40) and +7.8% AP (70-10). These gains, complemented by significant increases in AP₅₀, validate Q-MCMF’s ability to mitigate background foregrounding while optimizing the stability-plasticity balance in incremental object detection.

Stability-Plasticity Balance. Fig. 4 presents the performance of base and new classes under the 70-10 incremental setting for different α and β values. To study their individual effects, one parameter was varied while the other was fixed: β was fixed at 0.5 for α studies, and α was fixed at 0.7 for β studies. As shown in Fig. 4, optimal stability for base classes is observed at $\alpha = 0.7$ in (a) while optimal plasticity for new classes is at $\beta = 0.5$ in (b), with deviations in either direction causing performance deterioration. By observing new performance in Fig. 4 (a), new class performance consistently improves with increasing α . This occurs because elevated α reduces base class matching, thereby limiting in-

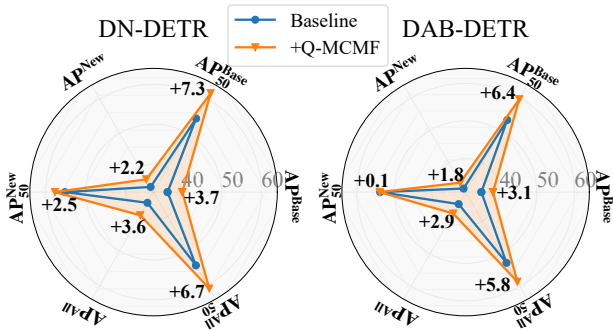


Figure 5: Effectiveness of Q-MCMF Matcher on DN-DETR and DAB-DETR under 70-10 setting.

terference from base to new class feature distributions. Symmetrically, in Fig. 4 (b), base class performance improves with higher β due to its equivalent protective effect against interference. The All curve shows that the overall performance represents a trade-off between stability and plasticity. This indicates the balance of stability and plasticity is manipulated by both α and β instead of α solely controls stability and β solely controls plasticity.

Effectiveness of Q-MCMF on Different Architectures.

To validate the robustness of our proposed method, we further evaluate our method on two extra architecture: DN-DETR and DAB-DETR. As shown in Fig. 5, we present the results of these two detectors under 70-10 setting. Specifically, our method demonstrates consistent improvements across both architectures. For DN-DETR, our method obtain a gain of 3.7% compared to our baseline while AP_{50} achieves a substantial +7.3% gain. New-class performance also increases, with AP increasing by +2.2% and AP_{50} by 2.5%. Consequently, overall AP is increased by +3.6% to 39.7%, and AP_{50} is increased by +6.7% to 60.9%. Similarly, DAB-DETR also exhibits noticeable gains. These consistent improvements across both DN-DETR and DAB-DETR architectures collectively demonstrate the robustness and generalizability of our approach, validating its effectiveness for transformer-based incremental object detection.

t-SNE Visualization. To further verify that our method causes less distortion to learned feature semantics compared to baselines, we perform qualitative analysis via t-SNE (Maaten and Hinton 2008) visualization. We random select 10 categories and extract the features of the decoder’s last layer and map them to 3D space using t-SNE. As shown in Fig. 6, features from our method exhibit stronger intra-class compactness (e.g., sheep and donut) and clearer inter-class separability (e.g., sheep and elephant), whereas baselines show more scattered clusters and overlapping between classes. This indicates that our method better preserves the original semantic structure of features, confirming its advantage in mitigating feature distortion and forgetting.

Match Visualization. Fig. 7 validates our motivation by contrasting Q-MCMF and Hungarian matcher decisions, showing how Q-MCMF eliminates geometrically implausible matches. Crucially, the Hungarian Matcher persists in assigning targets to geometrically incompatible predic-




Figure 6: t-SNE visualization of learned categories’ features after the second phase training of the 70-10 setting.

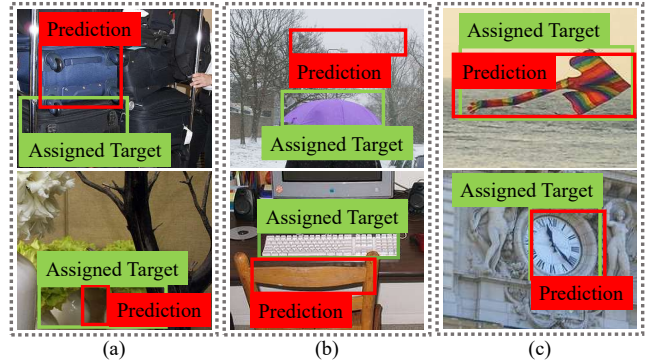


Figure 7: Zoomed in examples of matches. (a) shows examples of rejected match from base classes, (b) shows examples of rejected match from new classes, and (c) shows examples of accepted match.

tions for both base classes and new classes, causing background foregrounding. In contrast, Q-MCMF explicitly rejects such spatially inconsistent matches through quality-guided edge pruning, which significantly reduces background foregrounding and prevents erroneous supervision during training. More samples are presented in the appendix.

5 Conclusion

In transformer-based incremental object detection, we identify background foregrounding as a novel source of catastrophic forgetting. This phenomenon stems from misalignment between background predictions and foreground targets, directly caused by the Hungarian matcher’s exhaustive one-to-one matching requirement. To address this limitation, we propose the Q-MCMF matcher, which incorporates quality-guided edge pruning to eliminate geometrically implausible matches while simultaneously optimizing for minimal matching cost and maximal valid assignments. By preventing erroneous supervision from background foregrounding, our method effectively mitigates forgetting. Our work highlights the criticality of architecture-specific forgetting sources in IOD, with Q-MCMF offering a simple yet effective solution. We hope that our research will offer fundamental insights into forgetting mitigation for DETR-based incremental detectors, facilitating progress in this area.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62576282, 62176198, 62476223, 62576262, 62476226; in part by the National Key Research and Development Program of China under Grant 2024YFF1306501; in part by Innovation Capability Support Program of Shaanxi (Program No. 2024ZC-KJXX-043, 2024GX-YBXM-135); in part by Natural Science Basic Research Program of Shaanxi Province (2024JC-DXWT-07).

References

- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Cermelli, F.; Mancini, M.; Bulò, S. R.; Ricci, E.; and Caputo, B. 2020. Modeling the background for incremental learning in semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9233–9242.
- Chen, Q.; Chen, X.; Wang, J.; Zhang, S.; Yao, K.; Feng, H.; Han, J.; Ding, E.; Zeng, G.; and Wang, J. 2023. Group detr: Fast detr training with group-wise one-to-many assignment. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6633–6642.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Feng, T.; Wang, M.; and Yuan, H. 2022. Overcoming catastrophic forgetting in incremental object detection via elastic response distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9427–9436.
- Ge, Z.; Liu, S.; Li, Z.; Yoshie, O.; and Sun, J. 2021a. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 303–312.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021b. YOLOX: Exceeding YOLO Series in 2021. *arXiv preprint arXiv:2107.08430*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, S.; Lu, Z.; Cun, X.; Yu, Y.; Zhou, X.; and Shen, X. 2025. Deim: Detr with improved matching for fast convergence. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15162–15171.
- Jia, D.; Yuan, Y.; He, H.; Wu, X.; Yu, H.; Lin, W.; Sun, L.; Zhang, C.; and Hu, H. 2023. Detsr with hybrid matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19702–19712.
- Kim, J.; Cho, H.; Kim, J.; Tiruneh, Y. Y.; and Baek, S. 2024. Sddgr: Stable diffusion-based deep generative replay for class incremental object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28772–28781.
- Li, D.; Tasci, S.; Ghosh, S.; Zhu, J.; Zhang, J.; and Heck, L. 2019. RILOD: Near real-time incremental learning for object detection at the edge. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 113–126.
- Li, F.; Zhang, H.; Liu, S.; Guo, J.; Ni, L. M.; and Zhang, L. 2022. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13619–13627.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- Liu, S.; Li, F.; Zhang, H.; Yang, X.; Qi, X.; Su, H.; Zhu, J.; and Zhang, L. 2022. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*.
- Liu, Y.; Cong, Y.; Goswami, D.; Liu, X.; and Van De Weijer, J. 2023a. Augmented box replay: Overcoming foreground shift for incremental object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11367–11377.
- Liu, Y.; Schiele, B.; Vedaldi, A.; and Rupperecht, C. 2023b. Continual detection transformer for incremental object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23799–23808.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Luo, W.; Zhang, S.; Cheng, D.; Xing, Y.; Liang, G.; Wang, P.; and Zhang, Y. 2025. Gradient Decomposition and Alignment for Incremental Object Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4486–4495.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov): 2579–2605.
- Meng, D.; Chen, X.; Fan, Z.; Zeng, G.; Li, H.; Yuan, Y.; Sun, L.; and Wang, J. 2021. Conditional detr for fast training convergence. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3651–3660.
- Mo, Q.; Gao, Y.; Fu, S.; Yan, J.; Wu, A.; and Zheng, W.-S. 2024. Bridge past and future: Overcoming information asymmetry in incremental object detection. In *European Conference on Computer Vision*, 463–480. Springer.
- Peng, C.; Zhao, K.; Maksoud, S.; Li, M.; and Lovell, B. C. 2021. Sid: Incremental learning for anchor-free object detection via selective and inter-related distillation. *Computer vision and image understanding*, 210: 103229.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Shmelkov, K.; Schmid, C.; and Alahari, K. 2017. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, 3400–3409.

Wu, Q.; Zhang, S.; Cheng, D.; Xing, Y.; Xu, D.; WANG, P.; and Zhang, Y. 2025. Demystifying Catastrophic Forgetting in Two-Stage Incremental Object Detector. In *Forty-second International Conference on Machine Learning*.

Xing, Y.; Yang, S.; Wang, S.; Zhang, S.; Liang, G.; Zhang, X.; and Zhang, Y. 2024. MS-DETR: Multispectral pedestrian detection transformer with loosely coupled fusion and modality-balanced optimization. *IEEE Transactions on Intelligent Transportation Systems*.

Zhang, A.; Yang, D.; Liu, C.; Hong, X.; Shang, M.; and Zhou, Y. 2025a. DCA: dividing and conquering amnesia in incremental object detection. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence, AAAI'25/IAAI'25/EAAI'25*. AAAI Press. ISBN 978-1-57735-897-8.

Zhang, J.; Li, W.; Cheng, S.; Li, Y.; and Wang, S. 2025b. Dynamic Object Queries for Transformer-based Incremental Object Detection. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Zhang, S.; Chi, C.; Yao, Y.; Lei, Z.; and Li, S. Z. 2020. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9759–9768.

Zhang, S.; Lv, X.; Xing, Y.; Wu, Q.; Xu, D.; and Zhang, Y. 2025c. Revisiting Generative Replay for Class Incremental Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20340–20349.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.